REMARKS/ARGUMENTS

Claims 1-12 are pending. Claims 1, 4, 6, 7, and 11 have been amended to correct minor informalities. No new matter has been introduced.

Claims 7-9 stand rejected under 35 U.S.C. § 102(e) as being anticipated by Terrell et al. (US 2003/0189936A1).

Applicants respectfully submit that independent claim 7 is novel and patentable over Terrell et al. because, for instance, Terrell et al. does not teach or suggest a first storage device that is connected to the switch via the network; and a second storage device that is connected to the switch via the network. Nor does Terrell et al. disclose that if the data stored in the first storage device is stored in the second storage device, the switch converts a read request for the data stored in the first storage device into a data read request to the second storage device, and then transmits the converted data read request to the second storage device.

Terrell et al. discloses a router 102 having a managing process 204, a supervising process 206, and a plurality of routing processes 208. See paragraph [0084] and Fig. 2. The routing process 208 translates a virtual resource identifier included in a frame into a nonvirtual resource identifier to direct the frame to the nonvirtual resource through the port and network. This provides the data for a nonvirtual transaction frame (702). See paragraph [0161] and [0010]; and Fig. 7.

Terrell et al. does not disclose a first storage device (a physical entity) connected to the switch via the network. The virtual resource in Terrell et al. represents neither a physical entity nor a storage device connected to the router via a network.

Terrell et al. also fails to teach a second storage device connected to the switch via the network. Terrell et al. disclose a cache 424 that is contained in the router. See Fig. 4 and paragraph [0105]. The cache is not equivalent to the second storage device since no communication through a network of its own is involved by accessing the cache by the router. Further, a read request addressed to the cache is not an object of routing operation because the data handled by the managing processor does not involve routing process conducted by the routing processor. See Fig. 4. The managing processor in Terrell et al. has

a mirror agent process to be executed "by accessing cache 424, and directs port I/O process 402 to perform reads of the primary copy and writes the second copy to maintain the second copy." See paragraph [0119]. The second copy of data (the mirror) holds the same content as the primary copy of data. The mirror allows the user to obtain data from the second copy instead of the primary copy. The mirror function, however, is applied only to the cache 424, as described in paragraph [0119], not to a second storage device connected to the switch via the network.

Terrell et al. does not teach that if the data stored in the first storage device is stored in the second storage device, the switch converts a read request for the data stored in the first storage device into a data read request to the second storage device, and then transmits the converted data read request to the second storage device via the network. When accessing a cache instead of an remote storage, as in Terrell et al., a read request is not converted. Terrell et al. does not alter a resource identifier when accessing the cache, but merely translates a virtual resource identifier addressed to a virtual storage into a nonvirtual resource identifier.

The cache in Terrell et al. is not a second storage device as claimed, but a local storage that functions as a means for providing the user with the faster access data instead of accessing the data from the remote site more slowly. See paragraph [0118].

For at least the foregoing reasons, claim 7 and claims 8-9 depending therefrom are novel and patentable over Terrell et al.

Claims 1-6 and 10-12 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Terrell et al. in view of Moshfeghi et al. (US 6,779,119). The Examiner acknowledges that Terrell et al. does not disclose that the switch beforehand transfers data stored in the first storage device to the second storage device as recited in claims 1, 4, 6, and 11. The Examiner further acknowledges that Terrell et al. does not teach that if any amount of free storage capacity in the second storage device is not enough to store the data, the switch deletes some amount of data currently stored in the second storage device in a manner that data with the least frequency of use by the computer is deleted first, as recited in claim 10. The Examiner cites Moshfeghi et al. for supplying the missing teachings.

Applicants respectfully submit that independent claim 1 is patentable over Terrell et al. and Moshfeghi et al. because, for instance, they do not teach or suggest a first storage device that is connected to the switch via the network; a second storage device that is connected to the switch via the network; and a switch that beforehand transfers data stored in the first storage device to the second storage device via the network. Nor do they disclose that the switch converts the read request for the data stored in the first storage device into a data read request to the second storage device, and then transmits the converted data read request to the second storage device via the network.

Moshfeghi et al. discloses a system that prefetches information to hold it in a cache in anticipation of the user's request for reducing response time. See Abstract; and column 4, lines 38-56. The prefetcher 120 issues one or more anticipated commands through the task processor before the user submits a request corresponding to the command. The task processor obtains a response from the server via a network and the prefetcher 120 stores the response in a cache memory to prepare for the user's request.

Both Terrell et al. and Moshfeghi et al. disclose a cache (a local storage) that functions as a means for providing the user with the faster access data instead of accessing data in a remote site more slowly. The cache is not equivalent to a second storage device connected to the switch via the network. The transfer of data stored in the first storage device to the second storage device via the network is different from the transfer of data from the server to the local cache in Moshfeghi et al.

Moreover, as discussed above, Terrell et al. fails to teach or suggest that the switch converts the read request for the data stored in the first storage device into a data read request to the second storage device, and then transmits the converted data read request to the second storage device via the network. When accessing a cache instead of an remote storage, as in Terrell et al. and Moshfeghi et al., a read request is not converted.

For at least the foregoing reasons, claim 1 and claims 2-3 depending therefrom are patentable over Terrell et al. and Moshfeghi et al.

Applicants respectfully assert that independent claim 4 is patentable over Terrell et al. and Moshfeghi et al. because, for instance, they do not teach or suggest a first storage device that is connected to the switch via the network, and a second storage device

that is connected to the switch via the network, and a switch that beforehand transfers data stored in the first storage device to the second storage device. Nor do they disclose that the switch provides the computer with a third storage device corresponding to the first storage device, the third storage device being a virtual storage; and that the switch converts the data read request to the third storage device into a data read request to the second storage device, and then transmits the converted data read request to the second storage device via the network.

As discussed above, Terrell et al. and Moshfeghi et al. disclose a cache (a local storage) that functions as a means for providing the user with the faster access data instead of accessing data in a remote site more slowly. The cache is not equivalent to a second storage device connected to the switch via the network. Although Terrell et al discloses a virtual storage, it does not suggest that the virtual storage corresponds to a physical storage device.

In addition, both Terrell et al. and Moshfeghi et al. fail to teach that the switch converts the data read request to the third storage device (a virtual storage device) into a data read request to the second storage device, and then transmits the converted data read request to the second storage device via the network. When accessing a cache instead of an remote storage, as in Terrell et al. and Moshfeghi et al., a read request is not converted.

For at least the foregoing reasons, claim 4 and claim 5 depending therefrom are patentable over Terrell et al. and Moshfeghi et al.

Applicants respectfully contend that independent claim 6 is patentable over Terrell et al. and Moshfeghi et al. because, for instance, they do not teach or suggest a first storage device that is connected to the computer via a network; and a second storage device that is connected to the computer via the network; wherein the second storage device comprises a switch unit that is connected to the computer and the first storage device via the network, and a storage unit that is connected to the switch unit via an internal network. They further fail to disclose that when receiving the read request, the switch unit converts the read request for the data stored in the first storage device into a data read request to the storage unit, and then transmits the converted data read request to the storage unit.

As discussed above, Terrell et al. and Moshfeghi et al. disclose a cache (a local storage) that functions as a means for providing the user with the faster access data instead of accessing data in a remote site more slowly. The cache is not equivalent to a storage unit that is connected to a switch unit via an internal network. In addition, both Terrell et al. and Moshfeghi et al. fail to teach that the switch unit converts the read request for the data stored in the first storage device into a data read request to the storage unit, and then transmits the converted data read request to the storage unit via the internal network. When accessing a cache, as in Terrell et al. and Moshfeghi et al., a read request is not converted. Therefore, claim 6 is patentable.

Applicants respectfully assert that claim 10 is patentable over Terrell et al. and Moshfeghi et al. because, for instance, they do not teach or suggest a first storage device that is connected to the switch via the network; and a second storage device that is connected to the switch via the network, wherein if the data stored in the first storage device is stored in the second storage device, the switch converts a read request for the data stored in the first storage device into a data read request to the second storage device, and then transmits the converted data read request to the second storage device, as recited in claim 7 from which claim 10 depends.

As discussed above, Terrell et al. fails to teach the recited features. Moshfeghi et al. does not cure the deficiencies of Terrell et al., in that Moshfeghi et al. also discloses a cache (a local storage) that functions as a means for providing the user with the faster access data instead of accessing data in a remote site more slowly, instead of a second storage device. Moshfeghi et al. also fails to teach a switch that converts a read request for the data stored in the first storage device into a data read request to the second storage device, since when accessing a cache, as in Terrell et al. and Moshfeghi et al., a read request is not converted. Thus, claim 10 is patentable.

Applicants respectfully submit that independent claim 11 is patentable over Terrell et al. and Moshfeghi et al. because, for instance, they do not teach or suggest a switch having a port unit that is connected to an external device; a converter configured to convert commands and data which have been received by the port unit; and a switch unit configured to relay the command and the data according to address information. Nor do they disclose

that when receiving from the computer an access request for the data stored in the first storage device, the converter converts the access request into an access request to the second storage device, and that when receiving data corresponding to the access request from the second storage device, the converter converts the data into data transmitted from the first storage device, and then transfers the converted data to the computer.

As discussed above, Terrell et al. and Moshfeghi et al. disclose a cache (a local storage) that functions as a means for providing the user with the faster access data instead of accessing data in a remote site more slowly. The cache is not equivalent to a second storage device. Moreover, both Terrell et al. and Moshfeghi et al. fail to teach a converter that converts an access request for data stored in the first storage device into an access request to the second storage device, and that converts the data received from the second storage device into data transmitted from the first storage device. When accessing a cache instead of an remote storage, as in Terrell et al. and Moshfeghi et al., a read request is not converted. Nor is the data converted.

For at least the foregoing reasons, claim 11 and claim 12 depending therefrom are patentable over Terrell et al. and Moshfeghi et al.

CONCLUSION

In view of the foregoing, Applicants believe all claims now pending in this Application are in condition for allowance. The issuance of a formal Notice of Allowance at an early date is respectfully requested.

If the Examiner believes a telephone conference would expedite prosecution of this application, please telephone the undersigned at 650-326-2400.

Respectfully submitted,

Chun-Pok Leung
Reg. No. 41,405

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, Eighth Floor
San Francisco, California 94111-3834
Tel: 650-326-2400
Fax: 415-576-0300
RL:rl
60503492 v1